

Section 1: Introduction

This section provides a brief overview of IBM's Workload Manager, provides an overview of the WLM Component of CPExpert, describes the sources of data used by the WLM Component to analyze system performance, and describes the performance data bases that the WLM Component can use.

Chapter 1: IBM's Workload Manager

IBM's Workload Manager is a radical departure from earlier system management controls provided with MVS.

- With earlier versions of MVS, users provided MVS with detailed specifications about how MVS should process work.
 - Users classified workloads using the IEAICSxx member of SYS1.PARMLIB, and assigned the resulting workload classes to performance groups.
 - Domains and performance groups were defined in the IEAIPSxx member of SYS1.PARMLIB, and the domain and performance group specifications provided detailed guidance to MVS about how to process the work associated with the domains and performance groups.
 - System parameters contained in the IEAOPTxx member of SYS1.PARMLIB provided specifications to MVS at an overall system view. These specifications guided (or in some cases, actually controlled) MVS in how work was to be processed by the system. The specifications provided system-wide multiprogramming adjustment controls, system-wide expanded storage guidance (MVS/ESA SP4.2 introduced expanded storage guidance at the domain level), system-wide logical swap guidance, etc.

Systems personnel were required to understand the detailed internal logic of the System Resources Manager (SRM) part of MVS and how the system parameters effected that logic, in order to provide proper specifications in the SYS1.PARMLIB members. The MVS *Initialization and Tuning Guide* and *Initialization and Tuning Reference* documents described the logic and parameters, but the descriptions often were incomplete or misleading.

The SRM is complex, the concepts in the MVS *Initialization and Tuning Guide* and *Initialization and Tuning Reference* are difficult to appreciate, and SRM parameters often have unexpected results. Papers are regularly presented at professional conferences to describe newly-discovered performance problems or solutions resulting from experiences with the System Resources Manager¹. Further, many organizations

¹In fact, the whole concept of CPExpert arose because of a perceived need to have a tool to help analyze the performance of MVS.

did not even attempt to "tune" the SRM parameters because they did not understand the concepts. Many organizations simply used the IBM default values (even though IBM documents explicitly stated that the defaults should be changed).

- The Workload Manager attempts to solve these problem by removing the requirement that users provide detailed guidance to MVS on how MVS should process work. With the Workload Manager, users specify **performance goals**. The Workload Manager interprets the performance goals, and the Workload Manager tells MVS how to process work to meet the performance goals.

The major difference between earlier versions of MVS and the Workload Manager is that users do not provide the Workload Manager with detailed specifications on how to process the work assigned to a service class. Users describe **service** requirements, and the Workload Manager adjusts system resources to meet the service requirements.

IBM has provided a wealth of information about system performance and constraints to performance with the Workload Manager. The primary source of performance information about the performance of service classes is contained in SMF Type 72 (Subtype 3) records². The Type 72 records contain detailed information about performance goals, information which can be analyzed to determine whether the performance goals were met, and information showing performance constraints.

The Workload Manager is a new concept and new software. Consequently, there will be an industry-wide "learning curve" as installations gain experience using the Workload Manager to meet performance goals. The WLM Component of CPExpert can help reduce the learning curve. Additionally, the WLM Component will be improved as additional experience is gained by our users and by the industry as a whole.

Further, it is likely that some changes will be required in both the Workload Manager concept and the Workload Manager software to accommodate unique situations. The WLM Component of CPExpert will be modified as necessary to respond to these changes.

²The SMF Type 99 records contain extremely detailed information about the internal logic and decisions of the Workload Manager. The Workload Manager creates a "trace" of its logic as it examines service classes, makes resource adjustment decisions, makes policy adjustment decisions, makes working set management decisions, etc. The Type 99 records are written every 10 seconds, and may be extremely large. IBM strongly recommends that Type 99 records be collected for only a short time in response to specific needs. Since the Type 99 records are not normally collected, CPExpert does not consider them to be a **primary** source of performance data. In contrast, the Type 72 records are routinely collected at virtually every site.

Chapter 2: The WLM Component of CPEXpert

The WLM Component consists of numerous modules, working together to (1) detect intervals in which service classes did not achieve performance goals, (2) shape system performance and utilization data for detailed analysis by other modules, (3) evaluate the data to assess reasons service classes did not achieve performance goals, and (4) report the results from the evaluation. Additionally, the WLM Component contains modules to analyze the Workload Manager service definition to detect potential problems with workload classification or service class definitions. These modules are loaded and controlled by the central WLM Component of CPEXpert (titled WLMCPE).

- **Detect missed performance goals.** The WLM Component analyzes SMF Type 72 (Subtype 3) data to detect measurement intervals when service classes did not achieve response or execution velocity performance goals.
- **Shape Data.** The majority of "processing" in the WLM Component is accomplished by numerous modules whose function is to combine, sort, summarize, correlate, and prepare the data for analysis. These modules consolidate relevant data elements from SMF Type 70, Type 71, Type 72, Type 74, and Type 75 into one record for each service class that missed its performance goal during an entire RMF measurement interval. This record contains all information necessary to perform a comprehensive analysis of the system performance areas addressed by the WLM Component.

Additionally, data from SMF Type 30 records may be shaped and consolidated if the DASD Component of CPEXpert is licensed and the appropriate modification has been made to MXG or MICS code. Please refer to the DASD Component User Manual for details.

- **Evaluate Data.** The evaluation of the data is accomplished by rules whose purpose is to (1) evaluate the Workload Manager service policy, (2) identify problems or potential problems with the overall system, and (3) isolate reasons a service class did not meet its goal.
- **Report Results.** The WLM Component reports the results from the evaluation, in either "verbose" or "summary" reporting. The verbose reporting method presents a variety of data leading to a particular finding. The summary reporting method produces only a time-related graphical overview of the findings.
- **Maintain Historical Data.** The WLM Component optionally maintains a historical file of all rule results. This file is available for subsequent analysis by users who may wish to trend the frequency of different findings by the WLM Component.

Each rule that results in a finding is described in Appendix A of this document. The description summarizes the rule, lists predecessor rules, discusses the rationale for the finding, suggests action, and lists references related to the finding.

- The summary of the rule presents a short description of the finding.
- The predecessor rules are listed so you can follow the line of reasoning leading to a particular rule being executed.
- The discussion of the finding describes as much as necessary of the operation of the computer system (the Workload Manager, the SRM, MVS, the system hardware, etc.) as it relates to the particular rule. The purpose of the discussion is to explain the reasoning behind the rule, and what causes the rule to be produced. If appropriate, the discussion might refer you to related rules documented in the WLM Component User Manual.
- The suggestions list possible actions that should be considered based on the findings. In many cases, multiple possible actions are listed. You must determine which actions should be taken (this determination is based upon the suitability of the actions to your own environment, the financial implications of the action, and the "political" acceptability of the action.)
- References are listed with many rule descriptions. The references are provided so you can verify the analysis and suggestions.

Chapter 3: Data Sources

CPEXpert analyzes the performance of your system operating under the Workload Manager based upon data from two sources:

- **Guidance information.** Guidance information is contained in prefix.CPEXPERT.USOURCE(WLMGUIDE). The variables in the WLMGUIDE module can be viewed as "data selection and presentation" variables and "analysis control" variables.
 - The data selection and presentation variables allow you to select particular time intervals to be analyzed, and allow you to specify how the results from the analysis are to be presented.
 - The analysis control variables allow you to control the analysis the WLM Component will perform. These variables establish threshold values which are used by the WLM Component to assess system performance

The defaults provided for the analysis control variables will be appropriate for many environments. However, if your environment is unique, you can alter the analysis by changing the analysis control variables.

You should not hesitate to alter the guidance variables to meet your requirements. Some of the default values are deliberately set to cause rules to be produced initially. The purpose of this is to call your attention to the performance implications of certain decisions you may have made. Those decisions may be appropriate for your environment, even though CPEXpert may "flag" them as potential problems.

- **RMF information.** CPEXpert does most of its analysis based upon the information acquired by RMF and recorded into SMF Type 70(series) records. The primary information is contained in the SMF Type 72 records. IBM has included virtually everything necessary to analyze constraints to improved performance (from the view of the Workload Manager) in the Type 72 records: the service policy definition, the performance of service classes, delays to service classes, and performance and delays to subsystem transactions (if CICS Version 4 or IMS Version 5 are installed).

If installations do not operate RMF, they often operate a product that produces similar information. In particular, CPEXpert can analyze performance using the information provided the Capacity Management Facilities product from Boole & Babbage, Inc. if the data is available in a MXG or MICS performance data base.

Chapter 4: Performance Data Bases

The "raw" RMF data contained in the SMF Type 70(series) records must be translated into SAS format and placed into a SAS-based performance data base before CPEXpert can use the information. There are three ways in which these records normally are placed into a performance data base:

- **MXG performance data base.** The performance data base can be created by Merrill's Expanded Guide (MXG) software. MXG is provided by Merrill Consultants, Dallas, Texas. MXG provides a low-cost mechanism by which installations can create and maintain a performance data base.
- **MICS performance data base.** The performance data base can be created by the MVS Integrated Control System (MICS). MICS is provided by Legent Corporation, Vienna, Virginia. The MICS SMF (or Base) Component creates and maintains a comprehensive performance data base from SMF/RMF information.
- **Locally-created performance data base.** The performance data base can be one created locally from "raw" SMF Type 70(series) records. At present, no user of CPEXpert has applied CPEXpert to a locally-created performance data base. Please contact Computer Management Sciences for guidance if you wish to apply CPEXpert against a locally-created performance data base.

The flexibility to use any of the above performance data bases is due to the fact that CPEXpert is implemented in the SAS language. SAS provides a powerful macro facility, both with respect to macro coding and with respect to macro variable names.

CPEXpert uses SAS macro variable names when referring to an element of information in the SMF records. CPEXpert uses the SAS "%LET" statements to define the macro variables as MXG variable names, MICS variable names, or SMF basic variable names. These "%LET" statements are contained in unique variable definition modules for MXG, MICS, or raw SMF.

Thus, the same CPEXpert software can be executed against any of the three types of performance data bases, by only invoking the proper definition module. The SAS %LET statements in the definition module automatically cause CPEXpert to refer to the proper MXG, MICS, or SMF data elements.

Some important variables are not naturally available in the SMF Type 70(series) information, but are computed or derived by MXG and MICS. For these variables, CPEXpert arbitrarily chose the MXG variable name or chose an "artificial" SMF name [e.g., "SMF72CPU" represents the combined R723CCPU and R723CSRB times for service class periods represented in SMF Type 72 (Subtype 3) records].